GIANNIS DARAS

Assistant Professor, Massachusetts Institute of Technology

Cambridge, MA, USA • gdaras@mit.edu • giannisdaras.github.io • Google Scholar

Academic Appointments

Assistant Professor *July 2026 -*
MIT Sloan, Operations Research and Statistics
- Topics of interest: Generative AI, Data-centric AI, Inverse Problems, AI for Science.

Postdoctoral Associate *Nov 2024 - July 2026*
MIT, CSAIL
- Advisors: **Constantinos Daskalakis, Antonio Torralba**.

Education

Ph.D. Computer Science *August 2020 - Nov 2024*
The University of Texas at Austin
- Advisor: **Alexandros Dimakis**.
- Ph.D. Thesis: **Learning Generative Models from Corrupted Data.**

Electrical & Computer Engineering (5-year joint degree) *Sept. 2015 – June 2020*
National Technical University of Athens.
- *Overall GPA*: 9.24/10. Major: **Computer Science**. *Major GPA*: 9.5/10 (top 1%).

Selected Publications

1. **Giannis Daras***, Jeffrey Zhang*, Krithika Ravishankar, William Daspit, Costis Daskalakis, Qiang Liu, Adam Klivans, Daniel Diaz.
Spotlight, NeurIPS 2025. **Top 3.2%**.
Ambient Proteins: Training Diffusion Models on Low Quality Structures. [link], [code]

2. **Giannis Daras***, Adrian Munoz*, Adam Klivans, Antonio Torralba, Constantinos Daskalakis.
Spotlight, NeurIPS 2025. **Top 3.2%**.
Also, **Oral** at NeurIPS 2025 Workshop on Structured Probabilistic Inference.
Ambient Diffusion Omni: Training Good Models with Bad Data. [link], [code].

3. **Giannis Daras**, Alexandros Dimakis, Constantinos Daskalakis.
ICML 2024 and **Best Contribution Award at the International Biomedical and Astronomical Signal Processing (BASP)**.
Consistent Diffusion Meets Tweedie: Training Exact Ambient Diffusion Models with Noisy Data. [link], [code]

4. **Giannis Daras**, Kulin Shah, Aravind Gollakota, Alexandros Dimakis, Adam Klivans.
NeurIPS 2023.
Ambient Diffusion: Learning Clean Distributions from Corrupted Data. [link], [code]

Honors and Awards

- **Rising Star in AI**, Michigan, 2025.
- **Best Contribution Award** at the International Biomedical and Astronomical Signal Processing (BASP) Frontiers Conference 2025.
- **Spotlight, NeurIPS 2025** for the paper Ambient Omni (**Top 3.2%**).
- **Spotlight, NeurIPS 2025** for the paper Ambient Proteins (**Top 3.2%**).
- **Oral, NeurIPS 2023** for the paper Datacomp (**Top 0.58%**).
- **Oral, NeurIPS 2025 Workshop on Structured Probabilistic Inference & Generative Modeling** for the paper Multiresolution Textual Inversion.
- **Oral, NeurIPS 2025 Workshop on Structured Probabilistic Inference & Generative Modeling** for the paper Ambient Omni.

- **1st Award** at Fintech Hackathon by **National Bank of Greece**. (2019)
- **2nd Award** at "Be Finnovative 2.0 accelerator" by **National Bank of Greece**. (2019)
- **1st Place** at **"ECESCON9"** hackathon for our project BlindGuide. (2015)

Fellowships

- **UT Austin – Graduate Dean's Prestigious Supplement Fellowship.** *(2024)*
- **Onassis Ph.D. Fellowship** for the academic year 2023-2024.
- **Bodossaki Ph.D. Fellowship** for the academic year 2023-2024.
- **Leventis Ph.D. Fellowship** for the academic year 2023-2024.
- **Gerontelis Ph.D. Fellowship** for the academic year 2022-2023.
- **Hellenic Professional Society of Texas Ph.D. Award** (2022).
- **"The Great Moment of Education"** scholarship by **Eurobank EFG**. (2015)
- **Roumeliotis scholarship** for the highest grade in **Nationwide University Exam**. (2015)
- **ECE NTUA** scholarship for developing the official website of the department. (2017)

Teaching and Mentoring

MIT class 6.S982, Diffusion Models: From Theory to Practice Spring 2025
Designed and co-instructed a new diffusion class at MIT.
This graduate class attracted **84 students** and **113 listeners**.

UT Austin class EE381V: Advanced Machine Learning Spring 2023
Teaching Assistant – instructed 4 lectures. Evaluation from the students: 4.69/5.0.

Mentoring of over 15 students 2020-2025
Most mentoring experience has been acquired through MIT's Undergraduate Research Opportunity Program (UROP). Several students published their papers in top-tier ML conferences.

Invited Talks and Seminars

- **Rising Star in AI – Michigan AI Symposium** *(2025)*
- **Flatiron Institute (2025)**
- **Invited to be Core Participant** at the **Simons Institute** Semester Long Diffusion Program *(2027)*
- **MIT – Computer Vision Seminar** *(2025)*
- **Morgan Stanley Machine Learning Seminar Series** *(2025)*
- **MIT – ML Tea Seminar Talk** *(2025)*
- **Mila – Quebec AI Institute** *(2025)*
- **Ben-Gurion University** *(2025)*
- **Runway ML** *(2025)*
- **NTU Singapore** *(2025)*
- **Simons Institute Berkeley – The Future of Language Models and Transformers** *(2025)*
- **Columbia University – Workshop on Emerging Trends in AI** *(2025)*
- **Biomedical and Astronomical Signal Processing (BASP) Conference** *(2025)*
- **Harvard University, Sitan's Chen group meeting.** *(2024)*
- **Google DeepMind, London Office** *(2024)*
- **Grundfest Lecture series (UCLA + Caltech)** *(2024)*
- **Learning on Graphs and Geometry (LoGG) Reading Group** *(2024)*
- **Aalto University, Pervasive Computing Research Meeting** *(2024)*
- **University of Wisconsin-Madison, MLOPT Idea Seminar** *(2023)*
- **UT Austin – GenAI IFML Workshop** *(2023)*
- **EleutherAI Diffusion Reading Group** *(2023)*

- **Uppsala University, hosted by Jens Sjölund** *(2023)*
- **Archimedes Research Unit** *(2023)*
- **NeurIPS Workshop Oral Presentation, Multiresolution Textual Inversion** *(2022)*
- **Rice University, Imaging and Vision Seminar** *(2022)*

Media Coverage

The Independent article Discovering the Hidden Vocabulary of DALLE-2
Article Link

New York Post article Discovering the Hidden Vocabulary of DALLE-2
Article Link

Yahoo News Discovering the Hidden Vocabulary of DALLE-2
Article Link

Science Alert Discovering the Hidden Vocabulary of DALLE-2
Article Link

UT News AI Trained to Draw Inspiration From Images, Not Copy Them
Article Link

Institute of Foundations of Machine Learning (IFML) Ambient Diffusion
Article Link

National Science Foundation (NSF) Learning from noisy data listed as key direction for IFML.
Article Link

Work Experience

NVIDIA Research Scientist Intern *Oct 2023 - May 2024*
Research on diffusion models.

Archimedes – Research Visit *June 2023 - July 2023*
Research on learning generative models from corrupted data with Prof. Daskalakis.

Google – Student Researcher *February 2022 - July 2022*
Research on diffusion models.

Google – Research Intern *May 2021 - November 2021*
Research on solving inverse problems with GANs.

Explosion AI – Student Researcher *September 2018 – September 2019*
Research on Transformer architectures & sparse attention layers for spaCy.

Google Summer of Code *April 2018 – September 2018*
Integration of Greek language to spaCy and training of Greek Deep Learning models.

Service

Reviewing Service *2021 - Present*
NeurIPS (2021–2025), ICML (2021–2025), and ICLR (2023–2025), CVPR 2025, ECCV 2024, ICCV (2023, 2025), PNAS, TMLR, JMLR, Handbook of Numerical Analysis.

Graduate Admissions Students Committee, UT Austin *2022-2023*
Participated in the review of Ph.D. applications at UT Austin.

Workshop ReALM-GEN accepted at ICLR 2026 *2026*
Confirmed speakers: Volodymyr Kuleshon, Yong Chul Ye, Yaron Lipman.

Recommendation Letters

Alexandros Dimakis (Ph.D. advisor), **Constantinos Daskalakis** (MIT, Postdoc advisor), **Antonio Torralba** (MIT, Postdoc advisor), **Adam Klivans** (UT Austin)

Published Work

[Dar+20a] **Giannis Daras**, Nikita Kitaev, Augustus Odena, Alexandros G Dimakis. "Smyrf-efficient attention using asymmetric clustering". In: *Advances in Neural Information Processing Systems* 33 (2020), pp. 6476–6489.

[Dar+20b] **Giannis Daras**, Augustus Odena, Han Zhang, Alexandros G Dimakis. "Your local GAN: Designing two dimensional local attention mechanisms for generative models". In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2020, pp. 14531–14539.

[Dar+21b] **Giannis Daras**, Joseph Dean, Ajil Jalal, Alex Dimakis. "Intermediate Layer Optimization for Inverse Problems using Deep Generative Models". In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, 2021, pp. 2421–2432. URL: https://proceedings.mlr.press/v139/daras21a.html.

[Dar+22a] **Giannis Daras**, Negin Raoof, Zoi Gkalitsiou, Alex Dimakis. "Multitasking Models are Robust to Structural Failure: A Neural Model for Bilingual Cognitive Reserve". In: *Advances in Neural Information Processing Systems*. Ed. by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh. Vol. 35. Curran Associates, Inc., 2022, pp. 35130–35142. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/e45caa3d5273d105b8d045e748636957-Paper-Conference.pdf.

[Dar+22b] **Giannis Daras***, Yuval Dagan*, Alex Dimakis, Constantinos Daskalakis. "Score-Guided Intermediate Level Optimization: Fast Langevin Mixing for Inverse Problems". In: *Proceedings of the 39th International Conference on Machine Learning*. Ed. by Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato. Vol. 162. Proceedings of Machine Learning Research. PMLR, 2022, pp. 4722–4753. URL: https://proceedings.mlr.press/v162/daras22a.html.

[Dar+23a] **Giannis Daras**, Mauricio Delbracio, Hossein Talebi, Alex Dimakis, Peyman Milanfar. "Soft Diffusion: Score Matching with General Corruptions". In: *Transactions on Machine Learning Research* (2023). ISSN: 2835-8856. URL: https://openreview.net/forum?id=W98rebBxlQ.

[Dar+23b] **Giannis Daras**, Kulin Shah, Yuval Dagan, Aravind Gollakota, Alex Dimakis, Adam Klivans. "Ambient Diffusion: Learning Clean Distributions from Corrupted Data". In: *Advances in Neural Information Processing Systems*. Ed. by A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine. Vol. 36. Curran Associates, Inc., 2023, pp. 288–313.

[Dar+23c] **Giannis Daras***, Yuval Dagan*, Alex Dimakis, Constantinos Costis Daskalakis. "Consistent Diffusion Models: Mitigating Sampling Drift by Learning to be Consistent". In: *Thirty-seventh Conference on Neural Information Processing Systems*. 2023. URL: https://openreview.net/forum?id=GfZGdJHj27.

[Dar+24b] **Giannis Daras**, Weili Nie, Karsten Kreis, Alex Dimakis, Morteza Mardani, Nikola Borislavov Kovachki, Arash Vahdat. "Warped Diffusion: Solving Video Inverse Problems with Image Diffusion Models". In: *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. 2024. URL: https://openreview.net/forum?id=LH94zPv8cu.

[Dar+25a] **Giannis Daras***, Jeffrey Ouyang-Zhang*, Krithika Ravishankar, William Daspit, Costis Daskalakis, Qiang Liu, Adam Klivans, Daniel J. Diaz. "Ambient Proteins: Training Diffusion Models on Low Quality Structures". In: *Advances in Neural Information Processing Systems*. 2025. **(Spotlight)**.

[Dar+25b] **Giannis Daras***, Adrian Rodriguez-Munoz*, Adam Klivans, Antonio Torralba, Constantinos Daskalakis. "Ambient Diffusion Omni: Training Good Models with Bad Data". In: *Advances in Neural Information Processing Systems*. 2025. **(Spotlight)**.

[DCD25] **Giannis Daras***, Yeshwanth Cherapanamjeri*, Constantinos Costis Daskalakis. "How Much is a Noisy Image Worth? Data Scaling Laws for Ambient Diffusion." In: *The Thirteenth International Conference on Learning Representations*. 2025. URL: https://openreview.net/forum?id=qZwtPEw2qN.

[DDD24] **Giannis Daras**, Alexandros G. Dimakis, Constantinos Daskalakis. "Consistent diffusion meets tweedie: training exact ambient diffusion models with noisy data". In: *Proceedings of the 41st International Conference on Machine Learning*. ICML'24. Vienna, Austria: JMLR.org, 2024.

[Aal+25] Asad Aali*, **Giannis Daras***, Brett Levac, Sidharth Kumar, Alex Dimakis, Jon Tamir. "Ambient Diffusion Posterior Sampling: Solving Inverse Problems with Diffusion Models Trained on Corrupted Data". In: *The Thirteenth International Conference on Learning Representations*. 2025. URL: https://openreview.net/forum?id=qeXcMutEZY.

[Arv+22] Marius Arvinte, Ajil Jalal, **Giannis Daras**, Eric Price, Alex Dimakis, Jonathan I Tamir. "Single-shot adaptation using score-based models for mri reconstruction". In: *International Society for Magnetic Resonance in Medicine, Annual Meeting*. Vol. 42. 2022.

[CDD23] Sitan Chen, **Giannis Daras**, Alexandros G. Dimakis. "Restoration-degradation beyond linear diffusions: a non-asymptotic analysis for DDIM-type samplers". In: *Proceedings of the 40th International Conference on Machine Learning*. ICML'23. Honolulu, Hawaii, USA: JMLR.org, 2023.

[Gad+23] Samir Yitzhak Gadre, Gabriel Ilharco, Alex Fang, Jonathan Hayase, Georgios Smyrnis, Thao Nguyen, Ryan Marten, Mitchell Wortsman, Dhruba Ghosh, Jieyu Zhang, Eyal Orgad, Rahim Entezari, **Giannis Daras**, Sarah M Pratt, Vivek Ramanujan, Yonatan Bitton, Kalyani Marathe, Stephen Mussmann, Richard Vencu, Mehdi Cherti, Ranjay Krishna, Pang Wei Koh, Olga Saukh, Alexander Ratner, Shuran Song, Hannaneh Hajishirzi, Ali Farhadi, Romain Beaumont, Sewoong Oh, Alex Dimakis, Jenia Jitsev, Yair Carmon, Vaishaal Shankar, Ludwig Schmidt. "DataComp: In search of the next generation of multimodal datasets". In: *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. 2023. URL: https://openreview.net/forum?id=dVaWCDMBof. **(Oral Presentation)**.

[Jal+21] Ajil Jalal, Marius Arvinte, **Giannis Daras**, Eric Price, Alex Dimakis, Jonathan Tamir. "Robust Compressed Sensing MRI with Deep Generative Priors". In: *Advances in Neural Information Processing Systems*. Ed. by A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan. 2021. URL: https://openreview.net/forum?id=wHoIjrT6MMb.

[Li+24] Jeffrey Li, Alex Fang, Georgios Smyrnis, Maor Ivgi, Matt Jordan, Samir Yitzhak Gadre, Hritik Bansal, Etash Kumar Guha, Sedrick Keh, Kushal Arora, Saurabh Garg, Rui Xin, Niklas Muennighoff, Reinhard Heckel, Jean Mercat, Mayee F Chen, Suchin Gururangan, Mitchell Wortsman, Alon Albalak, Yonatan Bitton, Marianna Nezhurina, Amro Kamal Mohamed Abbas, Cheng-Yu Hsieh, Dhruba Ghosh, Joshua P Gardner, Maciej Kilian, Hanlin Zhang, Rulin Shao, Sarah M Pratt, Sunny Sanyal, Gabriel Ilharco, **Giannis Daras**, Kalyani Marathe, Aaron Gokaslan, Jieyu Zhang, Khyathi Chandu, Thao Nguyen, Igor Vasiljevic, Sham M. Kakade, Shuran Song, Sujay Sanghavi, Fartash Faghri, Sewoong Oh, Luke Zettlemoyer, Kyle Lo, Alaaeldin El-Nouby, Hadi Pouransari, Alexander T Toshev, Stephanie Wang, Dirk Groeneveld, Luca Soldaini, Pang Wei Koh, Jenia Jitsev, Thomas Kollar, Alex Dimakis, Yair Carmon, Achal Dave, Ludwig Schmidt, Vaishaal Shankar. "DataComp-LM: In search of the next generation of training sets for language models". In: *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. 2024. URL: https://openreview.net/forum?id=CNWdWn47IE.

[Rao+25] Negin Raoof, Litu Rout, **Giannis Daras**, Sujay Sanghavi, Constantine Caramanis, Sanjay Shakkottai, Alex Dimakis. "Infilling Score: A Pretraining Data Detection Algorithm for Large Language Models". In: *The Thirteenth International Conference on Learning Representations*. 2025. URL: https://openreview.net/forum?id=9QPH1YQCMn.

[Ren+24] Hui Ren, Joanna Materzynska, Rohit Gandikota, David Bau, **Giannis Daras**, Antonio Torralba. "Opt-In Art: Learning Art Styles Only from Few Examples". In: *Advances in Neural Information Processing Systems, Creative AI Track*. 2024.

[Rod+26] Adrian Rodríguez-Muñoz, William Daspit, Adam Klivans, Costis Daskalakis, Antonio Torralba, **Giannis Daras**. "Ambient Dataloops: Generative Models for Dataset Refinement". In: *Forty-third International Conference on Machine Learning*. 2026.

[Rou+23] Litu Rout, Negin Raoof, **Giannis Daras**, Constantine Caramanis, Alex Dimakis, Sanjay Shakkottai. "Solving Linear Inverse Problems Provably via Posterior Sampling with Latent Diffusion Models". In: *Thirty-seventh Conference on Neural Information Processing Systems*. 2023. URL: https://openreview.net/forum?id=XKBFdYwfRo.

[Sha+25] Kulin Shah, Alkis Kalavasis, Adam Klivans, **Giannis Daras**. "Does Generation Require Memorization? Creative Diffusion Models using Ambient Diffusion". In: *Forty-second International Conference on Machine Learning*. 2025. URL: https://openreview.net/forum?id=GGPM0z3dhU.

Preprints and Workshop Papers

[Dar+21a] **Giannis Daras**, Wen-Sheng Chu, Abhishek Kumar, Dmitry Lagun, Alexandros G Dimakis. "Solving inverse problems with nerfgans". In: *arXiv preprint arXiv:2112.09061* (2021).

[Dar+24a] **Giannis Daras**, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milanfar, Alexandros G Dimakis, Mauricio Delbracio. "A survey on diffusion models for inverse problems". In: *arXiv preprint arXiv:2410.00083* (2024).

[DD22a] **Giannis Daras**, Alex Dimakis. "Discovering the Hidden Vocabulary of DALLE-2". In: *NeurIPS 2022 Workshop on Score-Based Methods*. 2022. URL: https://openreview.net/forum?id=jxeSZaVzpmg.

[DD22b] **Giannis Daras**, Alex Dimakis. "Multiresolution Textual Inversion". In: *NeurIPS 2022 Workshop on Score-Based Methods*. 2022. URL: https://openreview.net/forum?id=3JCa_cqKaLy. **(Oral Presentation)**.

[Hos+25] Danial Hosseintabar, Fan Chen, **Giannis Daras**, Antonio Torralba, Constantinos Daskalakis. "DiffEM: Learning from Corrupted Data with Diffusion Models via Expectation Maximization". In: *arXiv preprint arXiv:2510.12691* (2025).

[Maj+26] Harris Abdul Majid, **Giannis Daras**, Francesco Tudisco, Steven McDonagh. *Ambient Physics: Training Neural PDE Solvers with Partial Observations*. 2026. arXiv: 2602.13873 [cs.AI]. URL: https://arxiv.org/abs/2602.13873.

[SDG26] Loukas Sfountouris, **Giannis Daras**, Paris Giampouras. *Align Invert: Solving Inverse Problems with Diffusion and Flow-based Models via Representation Alignment*. 2026. arXiv: 2511.16870 [cs.CV]. URL: https://arxiv.org/abs/2511.16870.

[Smy+23] Georgios Smyrnis, Matt Jordan, Ananya Uppal, **Giannis Daras**, Alex Dimakis. *Lovasz Theta Contrastive Learning*. 2023. URL: https://openreview.net/forum?id=-hWhz9xfrB9.

[Wei+26] Adam Wei, Nicholas Pfaff, Thomas Cohn, Arif Kerem Dayı, Constantinos Daskalakis, **Giannis Daras**, Russ Tedrake. *Ambient Diffusion Policy: Imitation Learning from Suboptimal Data in Robotics*. 2026. arXiv: 2606.12365 [cs.RO]. URL: https://arxiv.org/abs/2606.12365.